

WL-TR-97-1162



**THE AFFECT OF IMAGE COMPRESSION
ON A SYNTHETIC APERTURE RADAR
AUTOMATIC TARGET RECOGNITION
PRESCREENER AND THE RELATION TO
SAR IMAGE STATISTICS**

**RAYMOND WITHMAN
DONNIE CATES
BOB KOTZ**

**WRIGHT-LABORATORY
AUTOMATIC TARGET RECOGNITION DEV BRANCH
WL/AACR, 2241 AVIONICS CIRCLE
WRIGHT-PATTERSON AFB OH 45433**

AUGUST 1997

FINAL REPORT FOR PERIOD 01/95 - 06/97

Approved for public release; distribution unlimited

19980220 125

DTIC QUALITY INSPECTED 4

**AVIONICS DIRECTORATE
WRIGHT LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7623**

NOTICE

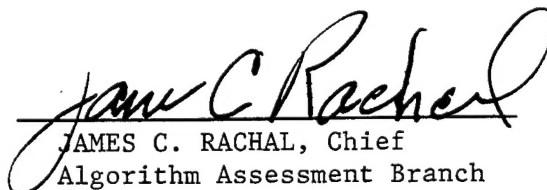
USING GOVERNMENT DRAWINGS, SPECIFICATIONS, OR OTHER DATA INCLUDED IN THIS DOCUMENT FOR ANY PURPOSE OTHER THAN GOVERNMENT PROCUREMENT DOES NOT IN ANY WAY OBLIGATE THE US GOVERNMENT. THE FACT THAT THE GOVERNMENT FORMULATED OR SUPPLIED THE DRAWINGS, SPECIFICATIONS, OR OTHER DATA DOES NOT LICENSE THE HOLDER OR ANY OTHER PERSON OR CORPORATION; OR CONVEY ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE, OR SELL ANY PATENTED INVENTION THAT MAY RELATE TO THEM.

THIS REPORT IS RELEASABLE TO THE NATIONAL TECHNICAL INFORMATION SERVICE (NTIS). AT NTIS, IT WILL BE AVAILABLE TO THE GENERAL PUBLIC, INCLUDING FOREIGN NATIONS.

THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION.



RAYMOND L. WITHMAN
Project Engineer



JAMES C. RACHAL, Chief
Algorithm Assessment Branch



JERRY L. COVERT, Chief
Combat Information Division
Avionics Directorate

IF YOUR ADDRESS HAS CHANGED, IF YOU WISH TO BE REMOVED FROM OUR MAILING LIST, OR IF THE ADDRESSEE IS NO LONGER EMPLOYED BY YOUR ORGANIZATION PLEASE NOTIFY WL/AACR WRIGHT-PATTERSON AFB OH 45433-7623 TO HELP MAINTAIN A CURRENT MAILING LIST.

Do not return copies of this report unless contractual obligations or notice on a specific document requires its return.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</p>			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	Final Jan 95 - Jun 97
4. TITLE AND SUBTITLE		5. FUNDING NUMBERS	
The Affect of Image Compression on a Synthetic Aperture Radar Automatic Target Recognition Prescreener and the Relation to SAR Image Statistics		PE 63226 PR 7629 TA 00 WU 00	
6. AUTHOR(S)		8. PERFORMING ORGANIZATION REPORT NUMBER	
Raymond Withman Donnie Cates Bob Kotz			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
Wright Laboratory Automatic Target Recognition Development Branch WL/AACR 2241 Avionics Circle WPAFB OH 45433		WL-TR-97-1162	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		12a. DISTRIBUTION AVAILABILITY STATEMENT	
Avionics Directorate Wright Laboratory Air Force Materiel Command Wright Patterson AFB OH 45433-7623 POC: Raymond L. Withman, WL/AACR (937) 255-1105 x3428		Approved for public release, distribution unlimited	
11. SUPPLEMENTARY NOTES		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words)		<p>High resolution SAR imagery compressed using three compression algorithms; Vector Quantization (VQ), JPEG and a Wavelet based MRES algorithm. The restored compressed imagery was processed through a morphology based prescreening Automatic Target Recognition algorithm and the results plotted on ROC curves. Based on these curves MRES and JPEG were found to perform significantly better than VQ. Several image statistics were computed for the restored compressed and uncompressed imagery. A prescreener performance figure of merit (FOM) was computed for each compression ratio for each algorithm. The correlation between each statistic and the FOM was computed. The results show the entropy is the statistic which correlates best with the FOM. Spearman's rank coefficient was used to perform a nonparametric analysis on the image statistics. The results confirmed the first results showing that entropy is the best predictor of prescreener performance. This analysis also showed that JPEG and MRES results are fundamentally different from VQ results. This is not unexpected since these two are based on transforms whereas VQ is a purely empirical technique.</p>	
14. SUBJECT TERMS		15. NUMBER OF PAGES	
Image Compression Synthetic Aperture Radar		Automatic Target Recognition SAR Image Statistics	
		45	
17. SECURITY CLASSIFICATION OF REPORT		18. SECURITY CLASSIFICATION OF THIS PAGE	
UNCLASSIFIED		UNCLASSIFIED	
19. SECURITY CLASSIFICATION OF ABSTRACT		20. LIMITATION OF ABSTRACT	
UNCLASSIFIED		UL	

Contents

1. Introduction.....	1
2. Compression Algorithms	2
2.1 JPEG.....	2
2.2 Vector Quantization.....	2
2.3 Multi-Resolution.....	3
3. Imagery.....	5
4. ATC Performance.....	7
4.1 ATC Concepts.....	7
4.2 Morphological Detector.....	10
4.3 Procedure.....	14
4.4 Discussion.....	14
5. Statistics of Compressed SAR Imagery.....	27
5.1 Statistics.....	27
5.2 Data Presentation.....	28
5.3 Tabular Data.....	28
5.4 Statistics Plots.....	29
5.5 Correlation Analysis.....	29
5.6 Rank Order Metric.....	29
5.7 Correlation Analysis.....	35
5.8 Performance Analysis.....	35
6. Entropy.....	38
7. Conclusions.....	41
8. Recommendations for Further Study.....	43
9. References.....	44

Preceding Page Blank

List of Figures

1. MBV ATR Paradigm.....	8
2. The ARAGTAP Morphological Detection Algorithm.....	11
3. Figure of Merit Definition.....	15
4. Prescreener ROC Curves without Fractal Discriminants.....	17
5. Prescreener ROC Curves with Fractal Discriminants.....	19
6. Comparison of Fractal vs. Non-Fractal Discriminant Performance.....	20
7. Figure of Merit without Fractal Discriminant.....	22
8. Figure of Merit with Fractal Discriminant.....	23
9. Algorithm Performance at 8:1 Compression Ratio.....	25
10. Performance Characteristics.....	26
11. Entropy and MNSD as a Function of Compression Ratio.....	33

List of Tables

1 Image Target Types.....	5
2 Imagery Compression Ratios.....	6
3 Performance Data.....	26
4 Statistics Summary.....	30
5 Statistics Rank Order.....	34
6 Rank Order by Class.....	37
7 Rank Order by Class.....	37
8 Spearman's Rank Correlation Coefficient.....	37

1. Introduction

Model based vision automatic target cueing and recognition systems processing synthetic aperture radar (SAR) imagery have advanced to the point where serious consideration can be given to incorporating them in operational systems. However, to date little work has been done on the impact of data compression on these MBV algorithms. It may be assumed that with today's technology neither of these items are really an issue. Although digital storage media capable of storing huge amounts of data are available, the need for compression still exists. Today's SAR systems are capable of collecting huge amounts of data, some of which may be stored for long periods of time. Storage requirements can typically be measured in terabytes. Although communications channels have very high bandwidths, there is no assurance that the channel will be available to a particular user during a period of high activity. Therefore, the less data a user has to transmit, the higher the probability that he will have access to the required channel.

Therefore, image compression is of interest for at least two reasons. First, image compression is used to reduce the bandwidth required of a communications channel used to transmit imagery from one location to another. Second, image compression is used to reduce the number of bits that are required to store the image for later analysis. It is anticipated that this imagery will be transmitted by means of communications channels that will require compression to ensure receipt in a timely manner. This paper addresses the impact of several compression algorithms on a specific MBV ATR prescreening algorithm. This paper further investigates the impact of various image compression algorithms on the statistics of synthetic aperture radar (SAR).

The affect of three compression algorithms on various SAR imagery statistics will be presented for three sets of imagery that have two resolutions. These statistics will be examined to determine if they can be used to predict the performance of an automatic target detection algorithm. Finally, this paper addresses the impact of the compression algorithms on the image entropy and de facto performance of synthetic aperture radar (SAR) imagery.

2. Compression Algorithms

Three compression algorithms were used in this effort each taking a different basic approach. The algorithms were the Joint Photographic Experts Group (JPEG)¹ compression algorithm, Vector Quantization² and a multi-resolution wavelet based sub-band³ approach. These particular algorithms were selected after a comprehensive survey of compression algorithms by ERIM.⁴ The major factors considered in this selection process included performance characteristics with SAR imagery and ease of implementation.

There are other justifications for selecting these particular algorithms. JPEG is an industry standard and is currently used in several DOD systems. Vector Quantization is currently being considered for another major DOD system, although in a different variation than that considered here. The multi-resolution approach using wavelet decomposition lends itself to progressive transmission that is currently receiving much attention. In progressive transmission, sub-bands are compressed and sent on a band by band basis. Each band is combined with previously sent bands to progressively build up the image quality of the receiving end. The advantage is that an operator at the receiving end can monitor image quality and terminate transmission when the quality is sufficient for his purpose thus eliminating the need to transmit the remaining sub bands. This further reduces the amount of data that needs to be transmitted over the reduction achieved in the compression algorithm.

2.1 JPEG

JPEG is basically a transform approach. It uses the discrete cosine transform (DCT) to compress the image^{5 6}. In our implementation, a two-dimensional DCT is applied to 8 x 8 pixel blocks. Additional compression is achieved by extracting the DC term from each image block. Since this term represents a constant value for each compression block, it can be compressed on a block-to-block basis using a differential pulse coded modulation (DPCM) technique.

2.2 Vector Quantization

Vector quantization techniques make use of the inherent vector characteristics of digital imagery⁷. Historically, compression algorithms, like DPCM, functioned by considering the value of each pixel independently. In VQ, groups of vectors are considered as one object to be processed. An image is divided into small chips, typically 8 or 16 pixels in a chip. They are arranged in a rectangular array, e.g. 2 x 4 or 4 x 4 pixels in a chip. Arranged so, each position

represents a dimension in an N-dimensional vector space. In our examples, N = 8 and 16., defined so each chip is a vector occupying some position in its N-dimensional space.

Compression is done by comparing each vector to a set of vectors which are representative of the actual vectors in the imagery. In order to achieve compression, the number of stored vectors has to be smaller than the possible values of the image vectors in the vector space. The stored vectors are encoded and stored as a "code book". As an example, a 2×4 , or eight-dimensional vector which consists of eight 8-bit pixels might be encoded into one 8-bit byte in the code book. This results in an 8:1 compression ratio. The performance of VQ obviously depends on how well the code book approximates the distribution and frequency of occurrence of its vectors in the image.

The VQ algorithm used by ERIM is a pruned tree structured VQ (TSVQ) algorithm^{8,9}. This algorithm employs a 4×4 pixel array to form a 16-dimensional vector. The TSVQ algorithm organizes the code words in a tree structure¹⁰, which eliminates the need for an exhaustive comparison of all of the code book entries. TSVQ reduces the amount of processing required from N operations to $\log_2(N)$, where N equals the numbers of entries in the code book. In the case where there are 256 entries in the code book, processing is reduced by a factor of $256/\log_2(256)$ or 32.

2.3 Multi-Resolution

This compression algorithm makes use of the wavelet transform to decompose the image into sub-images of varying frequency sub-bands. The wavelet transform is similar to the windowed Fourier transform. However, the wavelet transform optimizes the tradeoff between frequency resolution and spatial localization. The basics of wavelet theory are contained in "The Wavelet Handbook"¹¹. The wavelet transform makes use of a "mother" wavelet, which is then dilated and translated to yield the output. The wavelet transform results into a set of sparse coefficients, which allows for higher compression ratio. Since fewer non zero-coefficients are output from the transform, there is less data to be stored for image reconstruction. "The Wavelet transform provides an efficient method of decomposing a signal into sparse and meaningful coefficients that are highly related to the information content of the signal."¹² In this implementation, the 4 coefficient Daubechies filter^{13,14} is used. After the wavelet transform is done, the resulting coefficients are evaluated using a human visual system model. The number of bits used to encode a given coefficient is determined by the model's predicted response of the human to the corresponding frequency. Higher sensitivity coefficients are assigned more bits than

lower sensitivity coefficients. The number of bits assigned to a given coefficient is defined by equation 1.

$$B_k = B_{tot} + \frac{1}{2} \log \left\{ \frac{W_{HVS}(K)^2 / A_k}{\sigma_{gm}^2} \right\} \quad (1)$$

Where:

B_K = the average number of bits allocated to the detail image K,

B_{tot} = the overall average bit rate,

$W_{HVS}(K)$ = the human visual system weight obtained from the equations of Perkins and Lookabaugh¹⁵,

A_K = the relative area of the detail image K,

s_{gm}^2 = a weighted geometric mean of the squared $W_{HVS}(K)$ defined by equation 2:

$$\sigma_{gm}^2 = \frac{\left\{ \prod_{k=1}^m \left(\frac{W_{HVS}(k)^2}{A_k} \right)^{A_k} \right\}}{\sum_{k=1}^m A_k} \quad (2)$$

After bit allocation the coefficients were encoded using either a scalar or vector quantization depending on the number of bits assigned.

3. Imagery

The imagery used in this effort consisted of three sets of SAR imagery from two separate collections. One set of high resolution imagery, set A, and one set of low resolution imagery, set B, was collected in the Midwest in the fall of 1992. Ground truth from this collection is documented in the ground truth data report¹⁶. The second low resolution collection, known as set C, was made in the western US in 1985. Data set A consisted of 22 images, data set B consisted of 20 images, and data set C consisted of 8 images. Each image consisted of a 1024 x 1024 pixel image chip.

The imagery contained 13 different types of tactical targets as shown in Table 1.

Table 1: Image Target Types

Target Class	APC's	Tanks	Trucks	SP Howitzers	Anti-Aircraft Gun
Targets	BMP-1 BMP-2 BM-21 BTR-60 BTR-70 BTR-80	T-72 T-62	TZM Zil-131	2S1 2S3	ZSU-23/4

The images were processed by the Environmental Research Institute of Michigan (ERIM) as part of a SAR imagery data compression study. Processing consisted of compressing the imagery at various compression ratios for each algorithm. The images were then decompressed to restore the "original" image. The number of bits per pixel for the compressed imagery is given in Table 2. In this table, 8 bits/pixel represents the uncompressed imagery. Compression ratios as high as 64:1 were achieved for each algorithm.

Image analysts image truthed each image chip. Observed target locations were validated against ground truth and image pixel coordinates were derived for each valid target. These image truth locations were included as separate files associated with each image used in this effort. The image truth files allowed for automatic scoring of the tested algorithms.

Table 2: Imagery Compression Ratios

Bits /Pixel	8	4	2	1	.5	.25	.125
Set A							
JPEG	X	X	X	X	X	X	X
VQ	X			X	X	X	
MRES	X			X	X	X	
Set B							
JPEG	X	X	X	X	X	X	X
VQ	X			X	X	X	
MRES	X			X	X	X	
Set C							
JPEG	X	X	X	X	X	X	X
VQ	X			X	X	X	X
MRES	X	X	X	X	X	X	X

4. ATC Performance.

4.1 ATC Concepts.

4.1.1 General Concepts.

Before addressing the current problem it is helpful to discuss several concepts of image analysis and ATR. This is desirable to ensure that concepts are understood and that terms are defined.

The first sets of terms refer to the detail or degree to which an interpreter or automatic target recognizer can determine the nature of a target. Generally three levels of recognition are considered: detection, classification, and recognition . Detection involves determining that some sort of target exists within the scene. However, no additional information can be determined about the object. A target can be classified when the general type of the target can be determined. Classification can involve more than one level of class. Tracked vehicles, tanks, trucks, artillery, missiles, etc., are examples of target classes. A target is recognized when a tank could be determined to be a T-62 versus an M -1.

A distinction must be made between ATR and cueing. In this paper ATR will be used to refer to an automatic process that results in a final decision, either classification or identification, without human intervention. Cueing can have two meanings. In this report it will usually mean a process that results in information being presented to a human operator informing him of a probable target and some level of recognition. The human operator then intervenes to either make the final decision or validates the machine made decision. Cueing can also refer to the automatic tasking of a higher resolution sensor or a different mode of the same sensor when the ATR process makes a detection or classification. The purpose of this type of cueing can be to obtain a higher degree of recognition or to use the cue for target acquisition purposes. Obviously the optimal way to accomplish cueing is to do perfect identification.

4.1.2 Model Based Vision ATR.

4.1.2.1 MBV ATR Paradigm

Even though different degrees of recognition are desired from a particular machine, MBV-based ATR is progress serially through each level of recognition, increasing the level of knowledge about the object at each given level. The basic paradigm is illustrated by Figure 1. The three stages that are usually performed are detection, indexing, and identification. This paper addresses only the screening stage of a particular algorithm. Therefore, the discussion of

prescreening is more detailed, with the other parts of the paradigm receiving only a cursory treatment.

4.1.2.2 Screening

The screening stage processes every pixel in the image performing detection and discrimination to determine on a pixel-by-pixel basis whether each pixel is a likely target or is clutter. Two common ways of performing detection are Constant False Alarm Rate (CFAR) processing or simple thresholding. This is sometimes referred to as pixel detection, and these pixels are those used in further processing in an attempt to reduce false alarms. However, since ATRs are concerned with recognizing targets, a second stage of detection that can be called target detection is usually implemented. In this process, contiguous detected pixels are clustered to form target-sized detections. Often it is required that these detected targets have a minimum

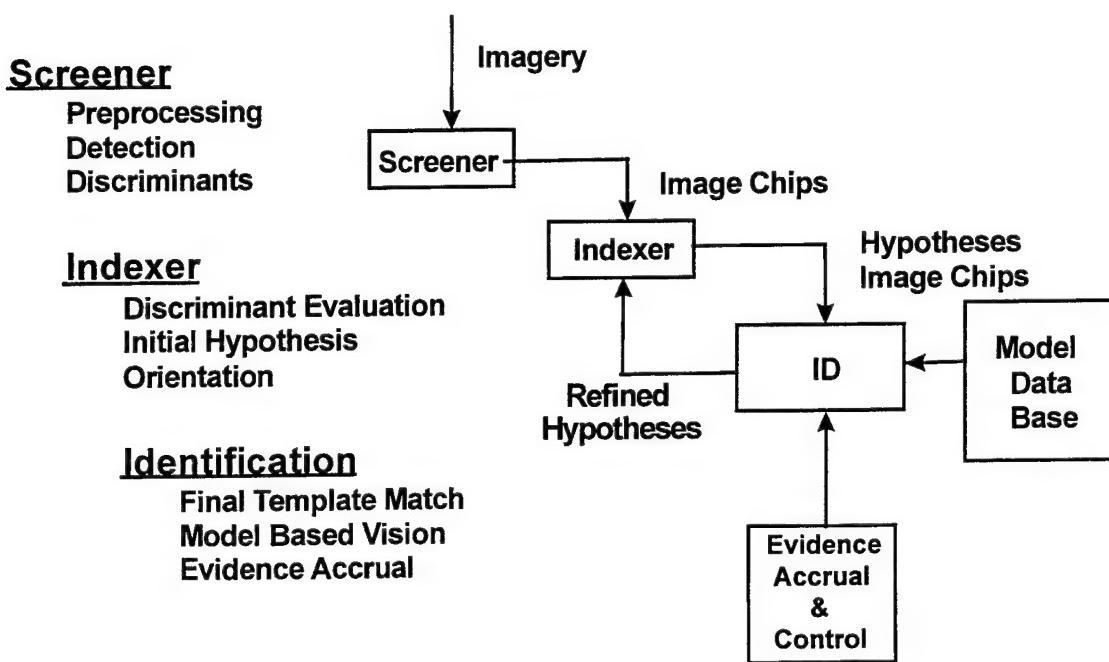


Figure 1. MBV ATR Paradigm

number of pixels to be considered a detection. This minimum size is dependent on the resolution of the radar and the types of targets that the ATR is designed to recognize.

In the screening stage, “target” detections are further processed to eliminate false detections. This screening stage can consist of several types of image processing. In statistical approaches, features are usually extracted and compared to values that represented the target types of interest. Typical features include aspect ratio, area, statistical measures, percentage of pixels detected in a bounding rectangle, etc. These features define dimensions in an N-dimensional feature, where N is the number of features that are extracted. Maximum likelihood decision logic is then used to discriminate between targets and non-targets.

Morphological processing is another approach that can be used to detect targets. Morphological processing differs from standard signal processing techniques which are usually linear operators based on statistical models assuming additive noise. Morphology is a set of non-linear operators based on non-parametric statistical methods assuming replacement type noise. It is effective at extracting shape and structure due in large part to an intuitive interpretation based on structuring elements that allow the experienced algorithm designer to efficiently feature extractor algorithms. Structuring elements^[17] are arrays which define the region of the image over which the morphological operation is to be performed. The structuring element is moved over the image and the operation is done for each pixel making up the image. The exact form of the structuring element determines what the output of a given morphological operation will be. Structuring elements are similar to the kernels of linear operations. Morphological processing can be done on either binary or gray scale images. Morphological processing can be used to estimate clutter, extract targets, smooth and fill in gaps in target signatures, and estimate target orientation or pose. Morphological processing also provides a simple technique for estimating fractal dimensionality. Fractal dimensionality has demonstrated promising results in separating man-made objects from natural clutter in SAR imagery. This paper considers a morphological detection algorithm with fractal dimensionality as a discriminant.

4.1.2.3 Indexing

Once targets are detected, the correct model must be selected for comparison to the imaged target. The process of selecting the most likely candidate models is called indexing. Target characteristics such as size, aspect ratio, and pose are used to determine which models to use. These characteristics allow the target to be classified into general classes, which limits the number of target models that have to be accessed.

4.1.2.4 Identification

Once likely targets are indexed, comparisons are made between the model and the candidate detection. These comparisons can consist of a large number of different techniques. A common technique is template matching using the complete target signature or only bright points.

Feature extraction and reasoning about feature spaces are also used. Various statistical techniques can also be employed. Another technique is to use external or collateral information to accrue knowledge about the return under test. In the ARAGTAP algorithm used in this study, this stage is referred to as the Prediction Description-Matching (PDM) loop.

4.2 Morphological Detector

The morphological detector used was developed for the ARAGTAP program¹⁸. It is more complex than the currently used CFAR detection in that it uses target shape as well as amplitude to perform the detection process. The algorithm is depicted in Figure 2. The boxes show the functions being performed at each stage of the algorithm. The notation at the side of each box indicates the morphological operation used to implement the function. Morphological operators used in this algorithm are discussed in reference¹⁹.

The imagery is first processed using a 3 x 3 close operator to remove speckle from the background. This operator tends to smooth the background while retaining the strong returns that should include the targets. Isolated high value returns will be removed, while more extended area of high return will be retained.

The next step is making an estimate of the clutter in the imagery. Once this is done, the estimated clutter "image" can be subtracted from the speckle reduced image to identify candidate target detections. This estimate is made using an N x N open operator. Making the N-dimension of the morphological structuring element sufficiently large ensures that large areas of contiguous high value pixels will be eliminated from the clutter estimate.

After the clutter estimate is made, the "clutter" image is subtracted from the "speckle" reduced image to find target candidates. Two thresholds are employed to provide two segmentations of the image. The "high" threshold is used to find very high valued pixels that are very likely parts of targets. This threshold is implemented so that a very low set percentage of the difference values is passed. Typically, less than .1 % of the pixels meet this criterion. Contiguous pixels that pass this threshold are grouped into "seeds" from which targets are grown in a later stage of the algorithm. A second "low" threshold is set to allow a large percentage of the difference values pass. Typically, 40% of the pixels pass this screen. This threshold results in groups of contiguous pixels that act to limit the candidate target size. These pixel groups are only used if the group contains at least one seed that passes the next processing step.

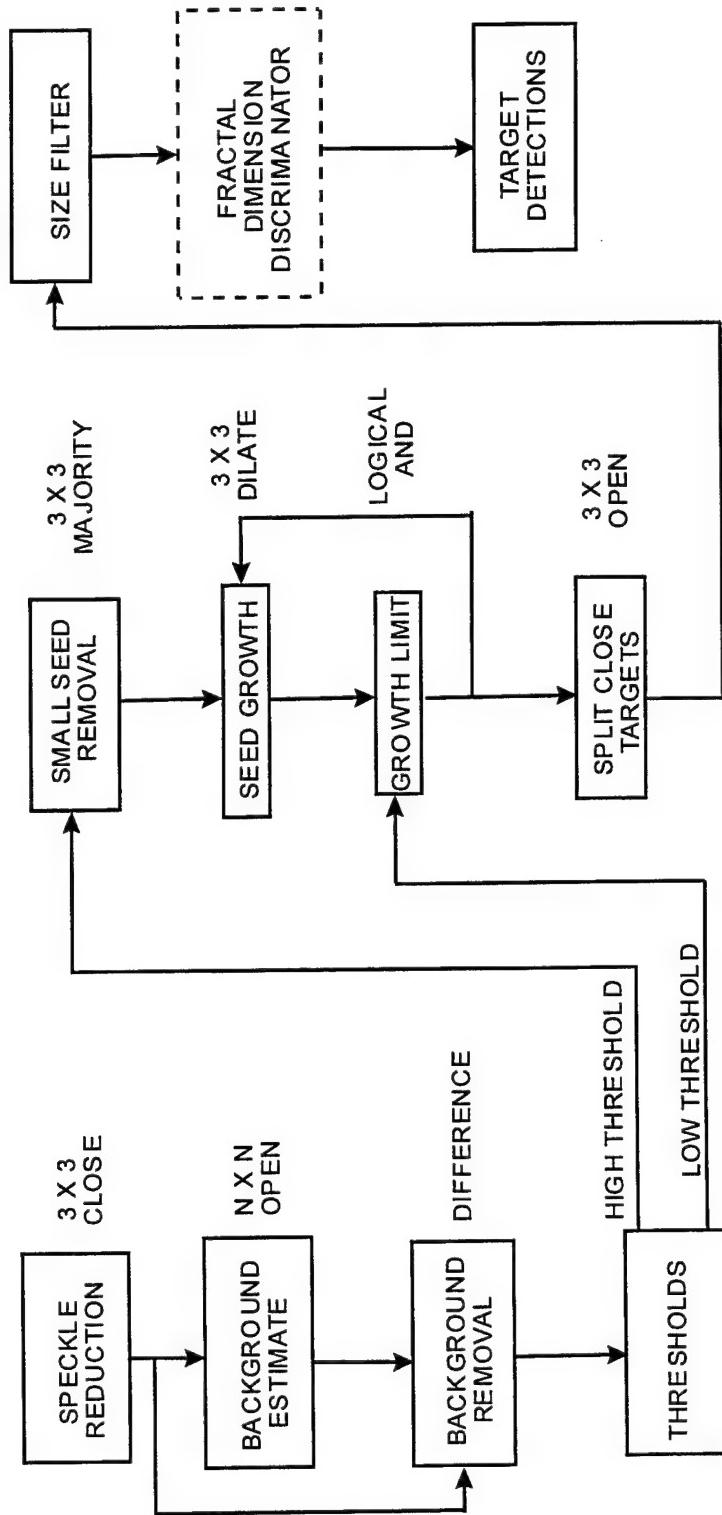


Figure 2 The RAGTAP morphological detection algorithm

After the clutter estimate is made, the "clutter" image is subtracted from the "speckle" reduced image to find target candidates. Two thresholds are employed to provide two segmentations of the image. The "high" threshold is used to find very high valued pixels that are very likely parts of targets. This threshold is implemented so that a very low set percentage of the difference values is passed. Typically, less than .1 % of the pixels meet this criterion. Contiguous pixels that pass this threshold are grouped into "seeds" from which targets are grown in a later stage of the algorithm. A second "low" threshold is set to allow a large percentage of the difference values pass. Typically, 40% of the pixels pass this screen. This threshold results in groups of contiguous pixels that act to limit the candidate target size. These pixel groups are only used if the group contains at least one seed that passes the next processing step.

The next stage is removing seeds that are judged to be too small to correspond to a real target. A 3 x 3 majority operator is used to do this. If the seed does not pass this screen, then it is eliminated from further consideration

Surviving seeds are assumed to be a part of a real target. However, since the "high" difference threshold allows only a minuscule number of these pixels to pass, it can be inferred that these pixels represent only the brightest scattering centers contained in the target. Therefore, these seeds are expanded or grown until they approximate in size actual targets of interest. This growth is done by using a 3 x 3 dilate operation on the seeds. This operator is applied successively until the seed has grown to the maximum allowed size. This maximum size is determined by the pixel arrays which have passed the "low" difference threshold. If the "high" threshold pixels can be thought of as the strong scattering centers, the "low" threshold pixels can be thought of as the weak scattering centers. Therefore, targets will generally be smaller than the pixel arrays which pass the "low" difference threshold. After each application of the dilate operator on the target seed, the result is compared to the "low" threshold array using a logical "and" operator. If the result is all true values, then the dilate is performed again. If a "false" value is found, then the seed has grown outside the limit defined by the "low" threshold array, and growth is terminated with the previous iteration's output being the expanded target.

The dilation operation used to grow the target seeds has the drawback of potentially merging two separate but closely spaced targets. This is remedied by performing a 3 x 3 close on the resulting expanded targets. This results in splitting closely spaced, weakly connected targets.

The final step in the detection process is a size filter. Target candidates that are too large or too small are eliminated from consideration. The remaining candidate targets are reported as detections.

4.2.1 Fractal Discriminator

In an additional processing stage, the fractal dimensionality is computed and used as a discriminator between man-made and natural objects. It could be argued that the fractal dimensionality more properly belongs as part of the classification process. However, in ARAGTAP it is considered part of the detection process. In this effort, detection was done using the algorithm with and without the fractal dimension discriminator as part of the detection algorithm.

Fractional dimensionality is computed using Richardson's Law (Eqn. 3), with the measured property being a morphological implementation of the covering blanket method^{20, 21,22}.

$$M(e) = K e^{d-D} \quad (3)$$

where

e = scale (=1,2,...)

M(e) = value of measured property at scale e

D = fractal dimension

K = constant

d = topological dimension plus one (=3 for gray scale image)

In this method, an image is processed using the dilate and erode operators. This results in two new surfaces. The dilation gives a surface that is above the original image, since each pixel is replaced by the maximum pixel value in the area of support. The erosion gives a surface that falls below the original image, since each pixel is replaced by the minimum pixel value in the area of support. The volume of the space between the eroded and dilated surfaces is then calculated. This volume is then converted into an area by dividing the volume by twice the current scale. This area is the measured property used to compute the fractal dimensionality. This process is repeated for several scales, and the results are used in Richardson's Law to determine the fractal dimension of the image. Results to date indicate that natural clutter has a constant fractal dimensionality at all scales, while the dimensionality of man-made objects varies monotonically as scale changes.

4.3 Procedure

To minimize the amount of processing it was desirable to make as few runs as possible. This meant varying as few parameters as possible and plotting as few points as possible. Preliminary runs were made to determine the performance sensitivity to algorithm parameters. It was determined that the most sensitive parameter was the high threshold. Other parameters were considered to have insufficient impact on the performance to be varied as part of this experiment. A second consideration was the number of points to be processed. It was determined that 10 points for each curve would be adequate to characterize the algorithm performance. Ten points adequately covered the knee of the curve for the non-fractal cases without generating a large number of points on the asymptotic part of the ROC. For the fractal cases some of the curves did not roll over the knee but were considered adequate since there was a point to point comparison of the performance between fractal and non-fractal versions of the algorithm. A script was used to process all the imagery for each compression algorithm, compression ratio. A scoring algorithm was written to make use of the image truth available for each image in the data set. The results were then plotted as described in the next section.

During this study no attempt was made to train the algorithm on compressed imagery. Training is desirable and would likely improve performance, i. e., minimize degradation.

4.4 Discussion

The results were plotted using two basic approaches. The first was the commonly used ROC curve. In this approach the probability of detection is usually plotted against the False Alarm Rate (FAR). We used a modified approach plotting the number of targets correctly detected versus the number of false alarms. Since the number of real targets and the total area covered is constant over the data set there is no information lost in this approach. Using this method allows us to present results that represent the degradation of performance that can be attributed to compression effects.

The second approach is based on an ATR users perspective. If a user requires a minimum P_D , then he is interested in how the false alarm performance is affected by the compression algorithms for the required P_D . This format is illustrated in Figure 3. In this method a figure of merit (FOM) is derived to indicate the false alarm effects of the various compression algorithms. The FOM is defined in the following manner. The detection rate is used as a parameter to generate the FOM. Several values for the detection rate are chosen and a perpendicular to the detection rate axis is drawn to intercept the ROC curves. At the points where

these intersections occur the number of false alarms are noted and normalized to the lowest FAR. These normalized values constitute the FOM. Thus we can determine false alarm performance as a function of compression ratio and/or compression algorithm given a constant detection rate. For this computation the detection rates were normalized by the highest detection rate found over all the ROC's. The normalized detection rates chosen were .7 and .8 since these values intersected the ROC curve for all situations examined. The FOM for detection rate of .9 was also plotted but as can be seen the number of points plotted is reduced in most instances since the performance dropped below the .9 level in many of the cases examined. The FOM's for detection rate of .8 are probably the most useful. At .7 the performance levels are still tightly clustered since the performance differences in processed imagery has not diverged enough to be significant. At .8, in most cases, the FOM points are on the lower portion of the knee of the curve and have diverged enough to show the effects clearly. This is also relevant since one tends to operate at the knee of the curve which is considered the best tradeoff between detection and false alarms in most applications. Note that there is no new information in this method. All the information is contained in the ROC curves.

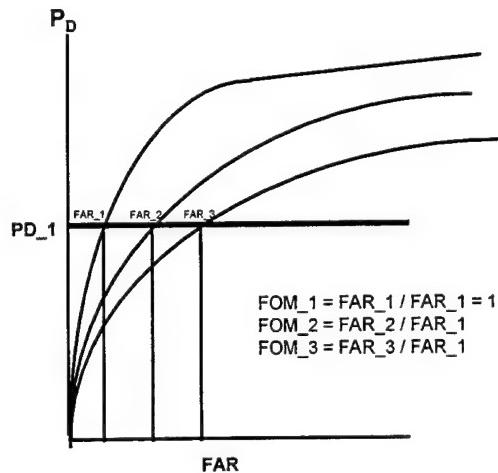


Figure of Merit Definition
Figure3: Figure of Merit Definition

4.4.1 Without Fractal Discriminator

Figure 4 depicts the ROC curves for the prescreener without the fractal discriminator. One of the characteristics of these curves is that the individual curves cross over each other. This is most pronounced for the VQ curves. As can be seen in Figure 2 the curve for 32:1 compression ratio is always better than the 8:1 and even crosses the uncompressed curve in the knee of the curve. Although most pronounced for VQ, similar crossings occur for JPEG and MRES. Since this occurs for all algorithms a large part of the affect must be attributed to detection algorithm characteristics. It could be inferred that the extreme effect seen in the VQ algorithm is caused by detection like behavior of the VQ algorithm. As the code book gets smaller only those vectors that resemble targets are saved. Therefore the compression process detects targets. This is a well-known property and other investigators have used it as the basis for automatic detection algorithms¹.

The JPEG algorithm shows a graceful degradation as the compression ratio increases. Although, there is no catastrophic breakdown in performance there does appear to be two breaks in the performance. Between 2:1 and 4:1 there is a gap in the ROC curves where performance deterioration accelerates. This is especially noticeable at the upper part of the knee extending into the asymptotic portion of the curve. A second gap occurs between 32:1 and 64:1. This gap is apparent in the lower part of the knee of the curve. Even so, performance is still fairly robust with detection declining from about 95% of peak performance to about 80% for 64:1 compression (about 87% for 32:1) in the knee of the curve.

The MRES results also show a graceful degradation. In general the degradation is smoother, but a small gap can be seen between 8:1 and 16:1 beginning at the upper end of the knee and extending through the rest of the curve. A shorter gap exists between 8:1 and 32:1 and is confined to the knee of the curve. Additionally, in the asymptotic part of the curve there is a crossover of the 32:1 curve resulting in very good performance at this compression ratio.

4.4.2 ROC Curves With Fractal Discriminator

Figure 5 shows the prescreener results using the fractal discriminator. In general the performance is better, as would be anticipated. The JPEG results show some crossovers but generally gracefully degrading performance. The gaps noticed before are mostly gone. However, a major performance falloff occurs between 32:1 and 64:1.

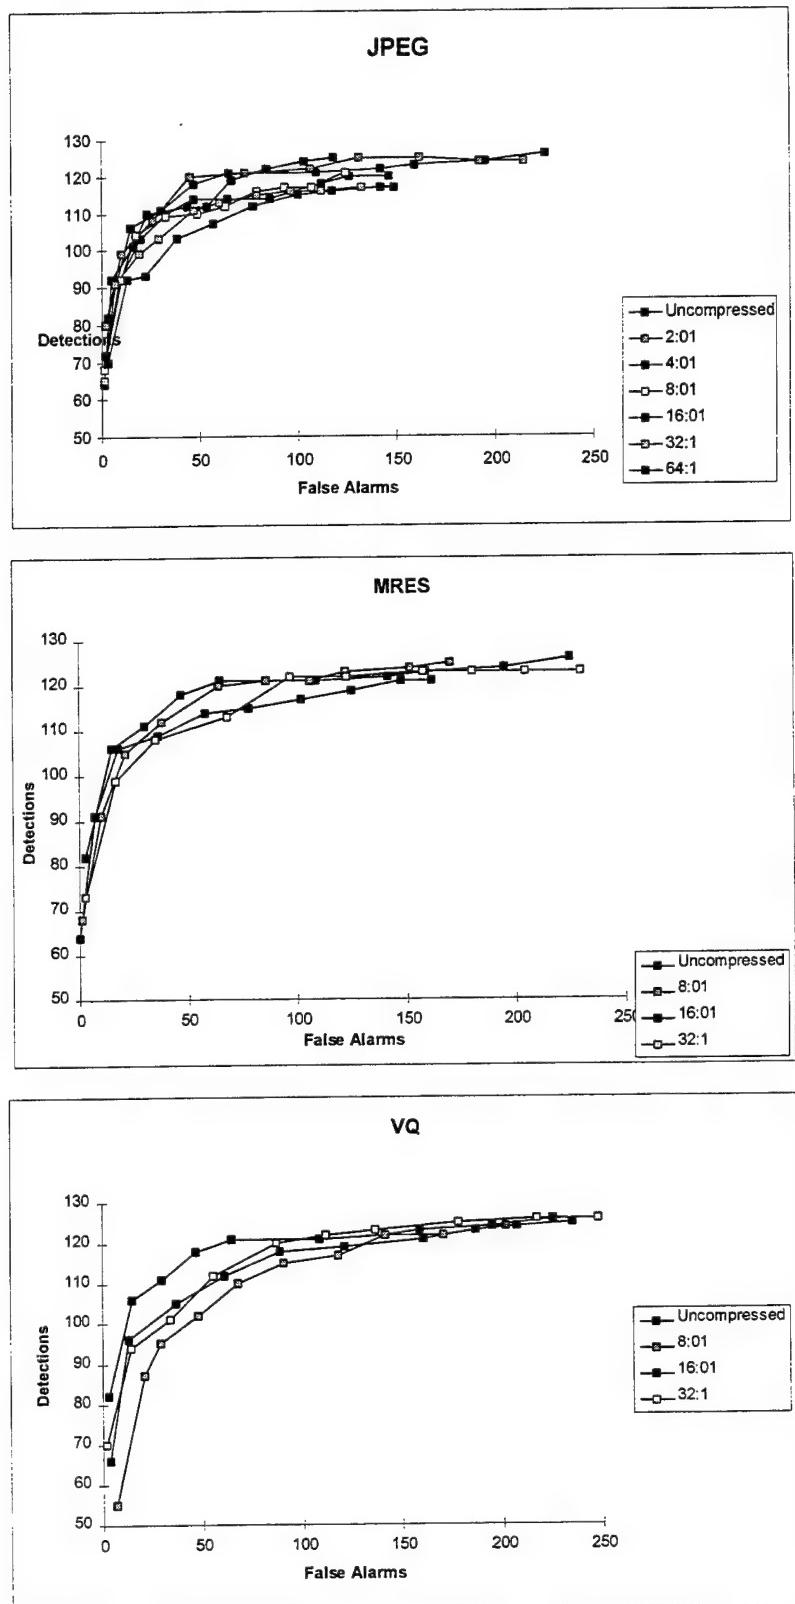


Figure 4. Prescreener ROC Curves without Fractal Discriminator

MRES performance generally degrades with increasing compression levels. However the gap between the 8:1 and 16:1 curves appears to widen and extend from the lower part of the knee throughout the rest of the curve. The uncompressed curve shows an instance where both the detection and false alarm performance both increased from point 4 to point 5. This does not occur in the non-fractal case.

VQ again shows the most erratic behavior. There is major divergence in performance at the lower end of the curve extending into the uppermost part of the knee. In the asymptotic portion of the curve all the curves indicate similar performance levels. Another anomalous situation is the increase in detection performance for two consecutive points in two sections of the curve with no corresponding increase in false alarms.

4.4.3 Fractal Discriminator

Figure 6 shows the impact of the fractal discriminator for uncompressed imagery and for 8:1 compressed imagery for all three compression algorithms. For uncompressed imagery the fractal discriminator reduces false alarms by a factor of nearly 10. Detections are reduced by about 10%. For JPEG, the fractal discriminator has reduced false alarms by about 3-4 to 1 and detections by about 10% compared to no fractal discriminator. For MRES the overall performance is improved by about 3 to 1 for false alarms with about a 10% reduction in detections. VQ results show that overall performance is improved by about 3 to 1 for false alarms with little decrease in detection performance

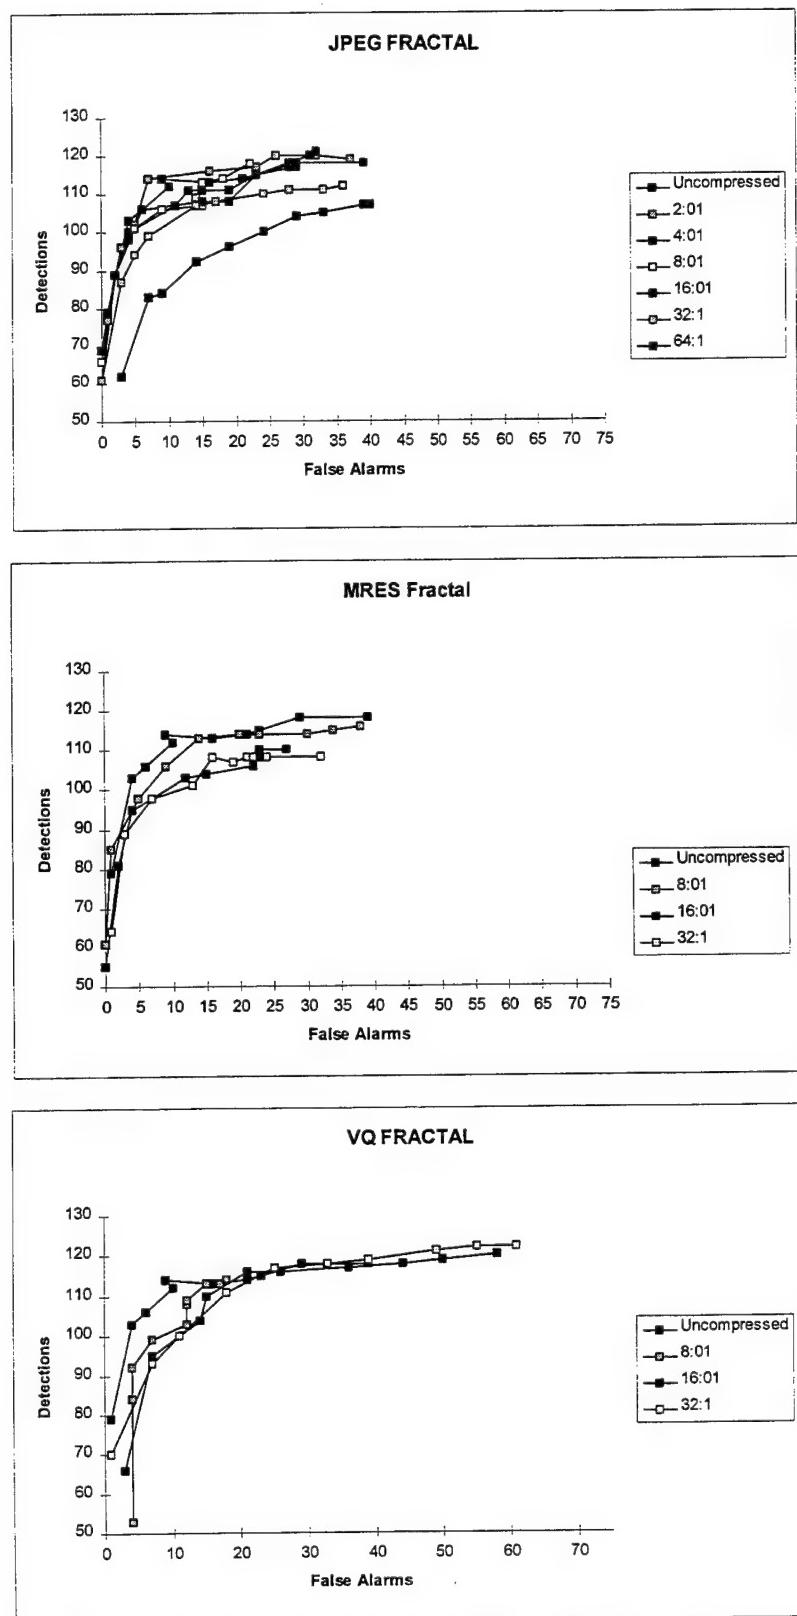
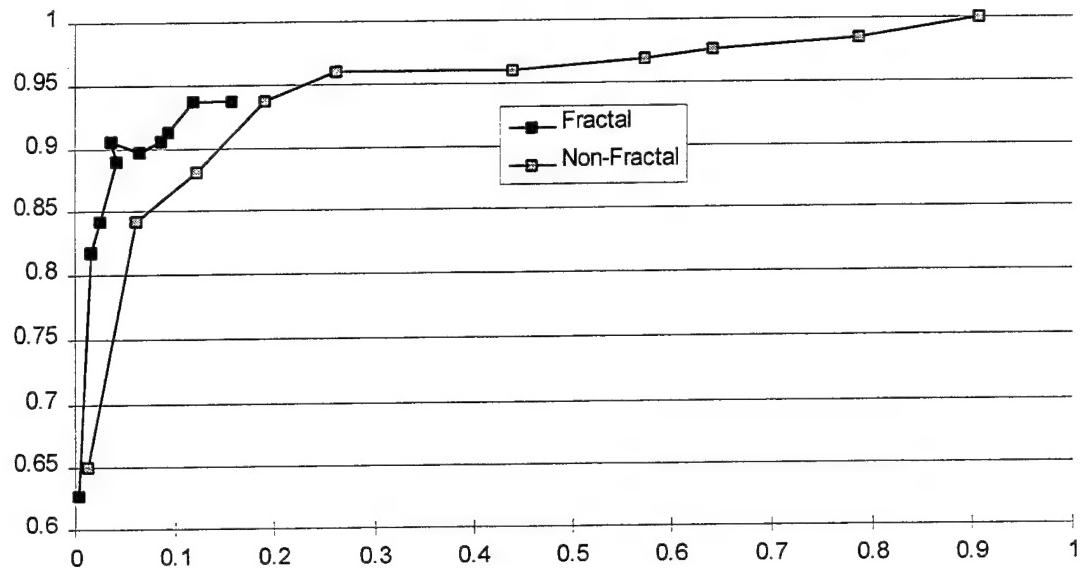


Figure 5. Prescreener ROC Curves with Fractal Discriminator

Uncompressed



FRACTAL 8:1 ROC CURVES

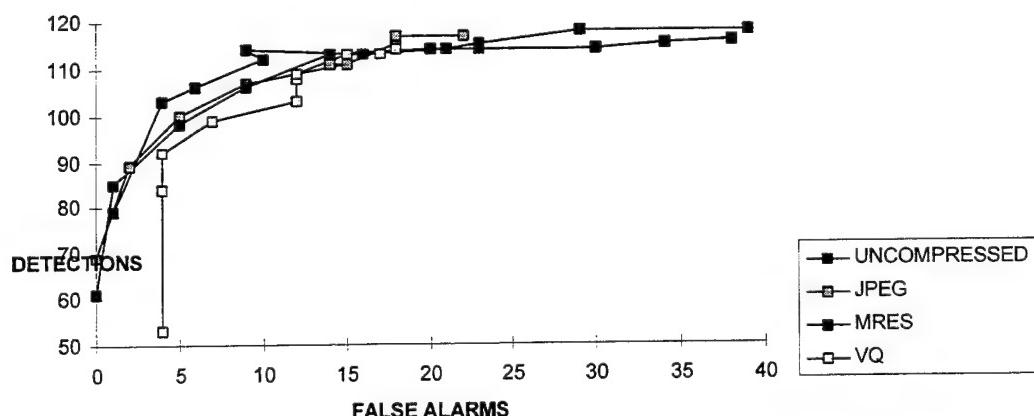


Figure 6. Comparison of Fractal vs. Non-fractal Discriminant Performance

Summarizing the impact of compression on the fractal discriminator, it can be said that compression affects detection performance very little. However, false alarms are significantly greater with compression than without, about a factor of three for 8:1 compression.

4.4.4 Figure of Merit (FOM)

The figure of merit defined in this paper is another way of looking at algorithm performance. It allows an evaluator to pick a specific detection performance level and see what affect the compression algorithms have on false alarms. A similar FOM could be defined by setting the false alarm rate and looking at how detection rates vary with compression ratios.

Figures 7 and 8 show FOM for the algorithm with and without the fractal discriminator respectively. These plots again show the relatively well behaved curves for the transform based algorithms and the erratic behavior of the VQ technique. Ideally one would expect smooth curves increasing monotonically as the compression ratio increases. Also, it would be expected that the curves would be ranked .7, .8., and .9 in order of increasing false alarms. Neither of these expectations hold. The curves generally do not fall in the proper order and often cross each other. This is just the result of the ROC curves crossing each other. Also, better performance by higher compression ratios accounts for the higher detection curves showing better performance than might be expected.

4.4.5 8:1 Compression Results

Since the results seem to indicate that compression can be done to 8:1 with little impact to detection performance, additional analysis was done on the 8:1 compression ratio results. Figure 9 shows the ROC curves for 8:1 compression performance. Figure 9a shows the algorithm without the fractal discriminator and figure 9b shows the results with the discriminator included. For the non-fractal case it can be seen that all the curves are well behaved and curve smoothly as expected. The VQ results, however are significantly lower than the other two algorithms. For the fractal case JPEG and MRES results are well behaved but the VQ shows the erratic behavior shown in the other charts. VQ performance is significantly worse in the lower part of the curve although false alarm performance at the high end seems to be better.

For the non-Fractal case MRES gives the best overall performance. JPEG shows similar performance at the low and far ends of the curve but is clearly not as good in the knee that is the most important part of the curve. VQ performance is significantly less than either MRES or JPEG. When the fractal discriminator is used similar results are obtained. VQ has inferior performance in

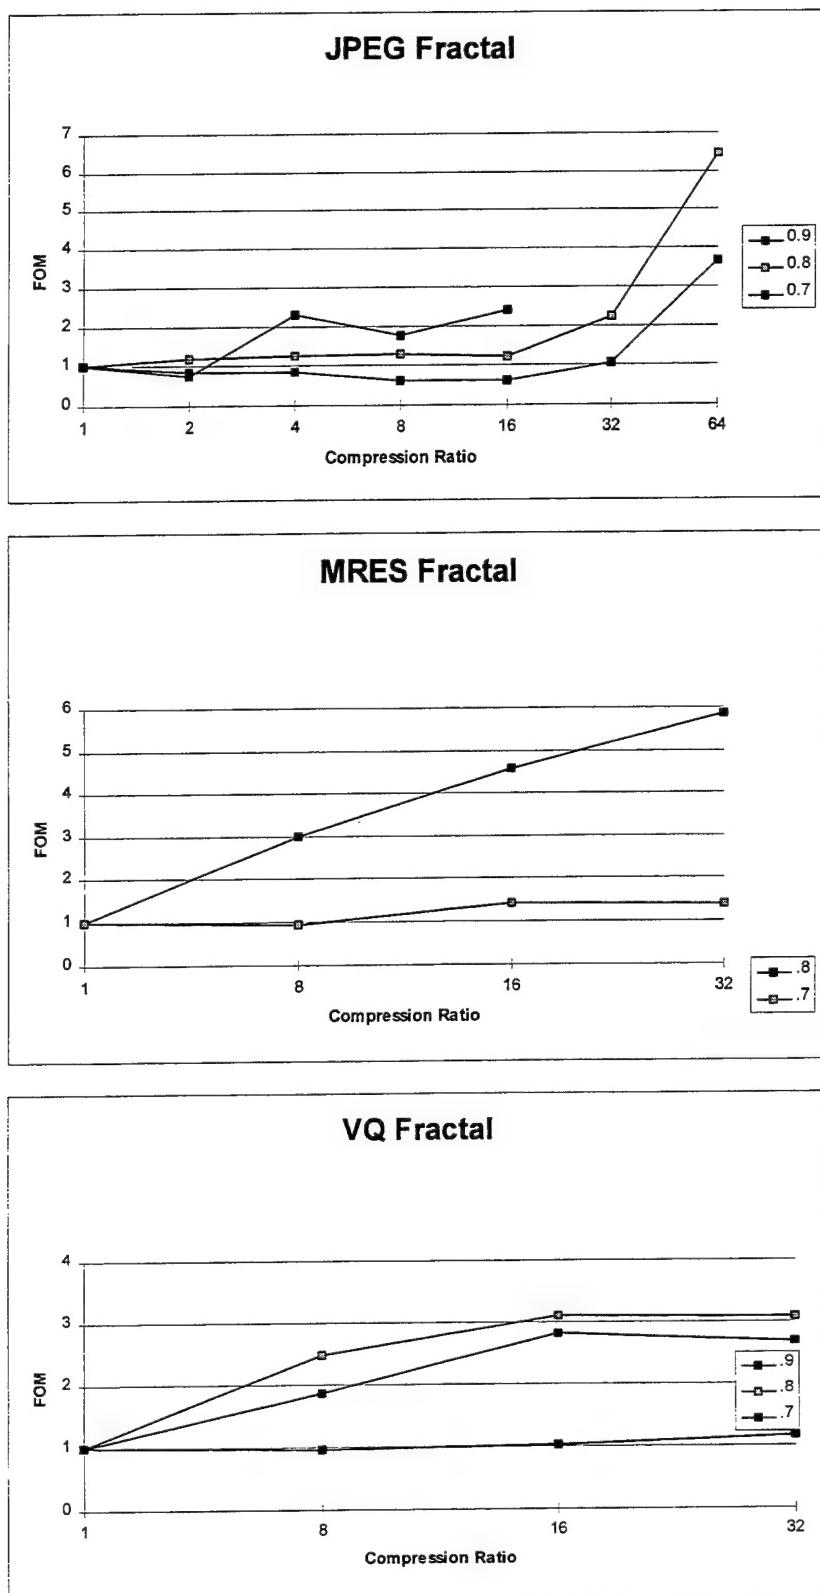


Figure 7. Figure of Merit without Fractal Discriminants

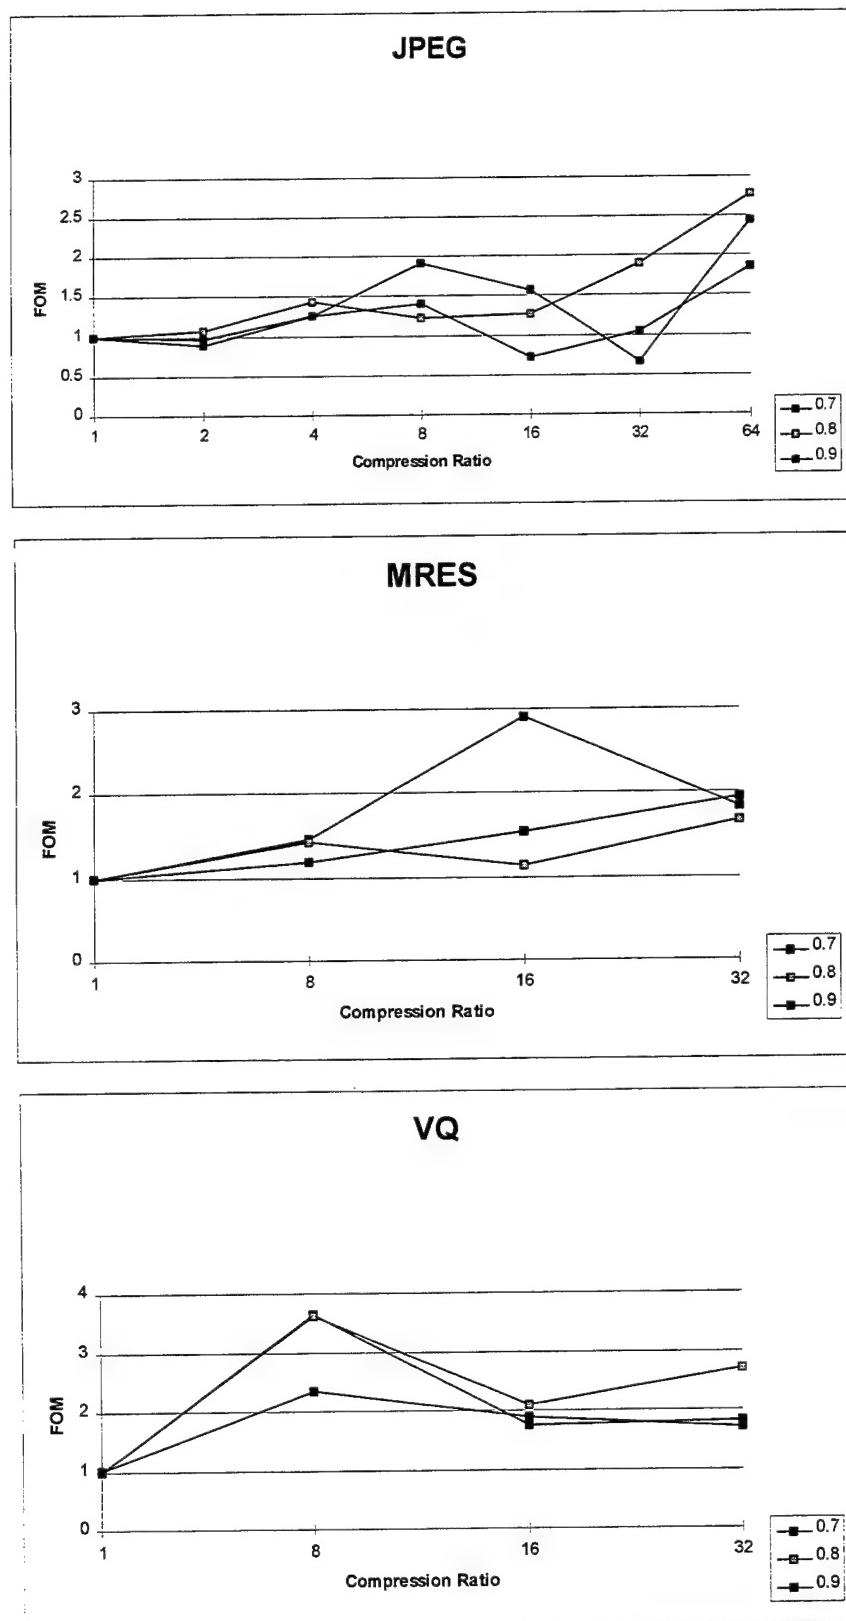
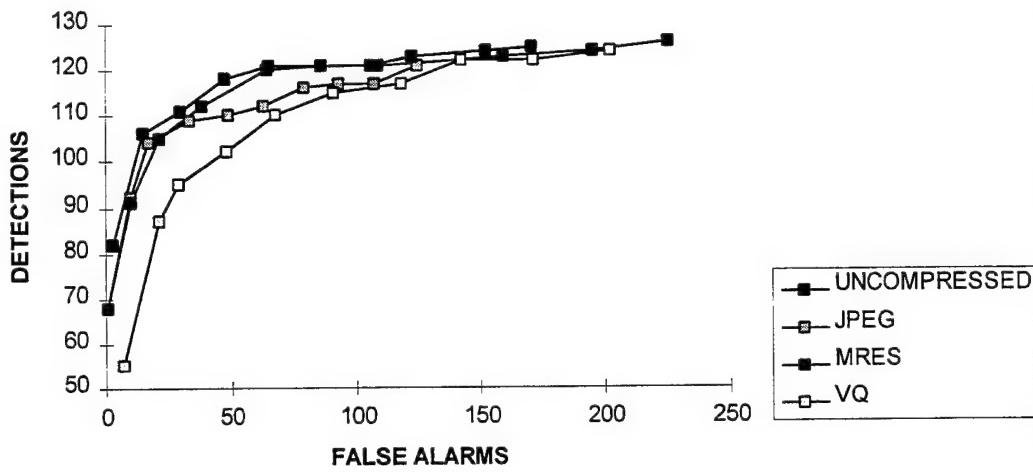


Figure 8. Figure of Merit with Fractal Discriminants

the knee of the curve and MRES and JPEG are nearly the same. Overall JPEG appears to perform a little better.

The ROC curves presented give a good qualitative or graphical view of detector performance. An attempt was made to reduce overall performance to a pair of numbers for each algorithm considered. The method used was to sum the detections and false alarms over the 10 values of the varied parameter (the high threshold). These results are plotted in figure 10. The points at the right are for the algorithm without the fractal discriminator. The fractal results are at the left. In both cases it appears that JPEG yields the best results. The percentage change in performance from the uncompressed results was calculated and is shown in table 2. This measure is very crude and probably distorts the actual overall performance, since it is not consistent with the qualitative evaluation done above. On the other hand this metric may say something about the relative robustness or sensitivity of the detection algorithm over a range of detection parameter settings.

NON-FRACTAL 8:1 ROC CURVES



FRACTAL 8:1 ROC CURVES

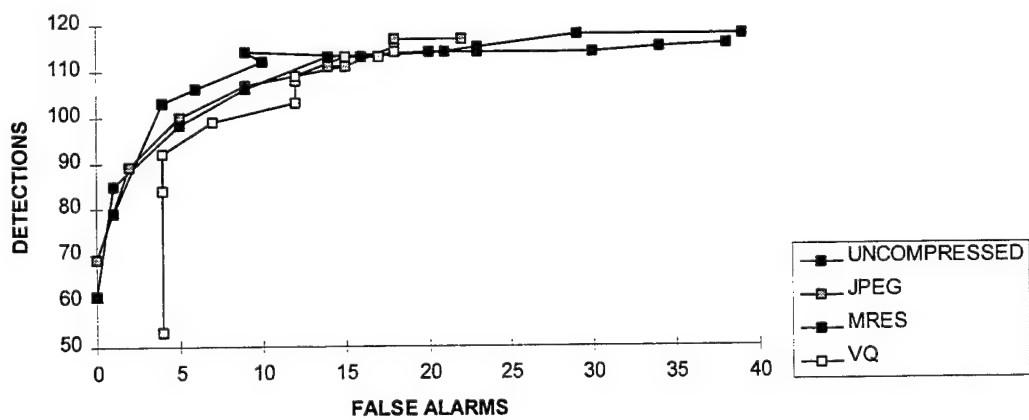


Figure 9. Algorithm Performance at 8:1 Compression Ratio

Performance Characteristics

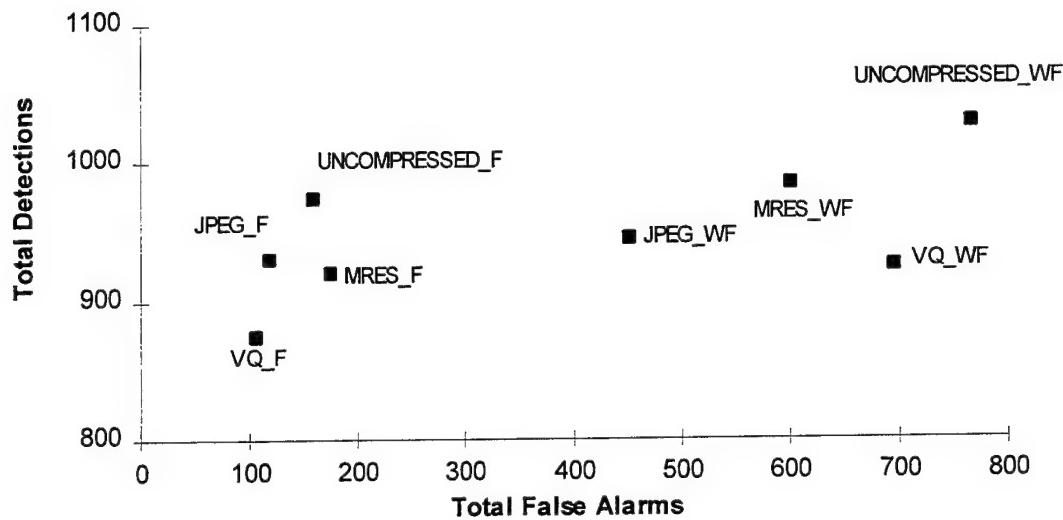


Figure 10. Performance Characteristics

Table 3: Performance Data

	FA's	Detections
W/O FRACTAL		
JPEG	-40.9	-8.1
MRES	-21.6	-4.2
VQ	-9.2	-10.0
WITH FRACTAL		
JPEG	-25.3	-4.5
MRES	+10.1	-5.4
VQ	-33.5	-10.3

5. Statistics of Compressed SAR Imagery

5.1 Statistics

Eight different statistics were measured for this study including entropy. These statistics are:

1. Entropy(ε):

$$\varepsilon = - \sum_{i=1}^n p(x_i) \log_2(p(x_i))$$

2. Peak (π):

$$\pi = \max\{x_1, x_2, x_3, \dots, x_n\}$$

the maximum pixel value in the image.

3. Mean (μ) - the first moment about zero

$$\mu = \frac{\sum_{i=1}^n x_i}{n}$$

4. Standard deviation (σ) - the square root of the second moment about the mean:

$$\sigma = \sqrt{\frac{\sum_{i=1}^n (x_i - \mu)^2}{(n-1)}}$$

5. Kurtosis (κ) - the fourth moment about the mean:

$$\kappa = \frac{m_4}{\sigma^4}$$

where m_4 is defined as;

$$m_4 = \frac{\sum_{i=1}^n (x_i - \mu)^4}{n-1}$$

Kurtosis measures the relative peakiness or flatness of a distribution.

$$6. \text{ Contrast1 (K1): } K1 = \frac{\sum_{i=1}^n x_i^2}{n}$$

$$7. \text{ Contrast2 (K2): } K2 = \frac{\pi}{\mu}$$

$$8. \text{ The mean normalized standard deviation (v) (MNSD): } v = \frac{\sigma}{\mu}$$

The first six statistics were computed using the Khoros descriptive statistics algorithms²³. The last two statistics were computed from the mean, standard deviation and peak statistics. A more standard contrast measurement

$$\text{contrast} = \frac{\max - \min}{\max + \min}$$

was considered but could not be used since nearly all minimum pixel values were 0 that would have resulted in all the values for contrast being 1 or nearly 1.

5.2 Data Presentation.

Data in this section is presented in three formats. The first is simply listing the data in a table. This is done as a way to present a comprehensive summary of the basic data in a relatively compact manner. The second format is plotting the values of the statistics as a function of the compression ratio. The third method is plotting the FOM as a function of compression ratio as described above

5.3 Tabular Data

Table 4 contains a summary of all the statistical data processed. The numbers represent the average value of the statistic for a given data set. Table 4a contains the results for data set A, table 4b the results for data set B, and table 4c the results for data set C. The first row of numbers indicates the compression ratio for each corresponding column of data. Examination of these tables shows that some of these statistics seem to vary as a function of the compression ratio, while others have no relationship to the compression ratio. For example ϵ and v seem to show a functional relationship while π and $K1$ seem to be unrelated to the compression ratio.

5.4 Statistics Plots

Using the information derived from the tables, ε and ν are plotted for data set A. Data set A was used because there are detection performance data available for this set. Statistics e and n are used since they seem to be the most likely to correlate automatic detection performance. The plots are contained in Figure 11. The Entropy metric is somewhat more stable as a function of compression ratio than the MNSD metric. However, both generally decrease as the compression ratio increases.

5.5 Correlation Analysis

Automatic target detection is expected to decrease with higher compression ratios as do the two metrics plotted in Figure 4. These two metrics were plotted along with the FOM for two parameter values designated VALUE1 and VALUE2. Only JPEG and MRES results were plotted since there appeared to be a weak correlation between the FOM and statistics for the VQ data. These plots are shown in Figure 3. To test the relationship among the detection performance and the behavior of image statistics the correlation measure was calculated for each set of data. The results were ranked by strength of the correlation. These results are shown in Table 4.

The results seem to indicate that entropy and the MNSD metric are the best indicators of detection performance for JPEG and MRES algorithms. Neither of these statistics ranked lower than third for these compression algorithms. For VQ compression the results are very different. There appears to be no consistency with the other algorithms. For VALUE 1 entropy and MNSD are ranked in the middle of the statistics. For VALUE2 MNSD is first and entropy is last. The most consistent statistics are peak and kurtosis, which are not highly ranked for

JPEG and MRES. For both values peak is significantly better than kurtosis.

5.6 Rank Order Metric

A rank order analysis was done on the data to determine the validity of the observations in the preceding paragraph. The six sets of data were grouped in six different arrangements and a rank order metric calculated. The metric is determined in the following manner. The rank of a particular statistic for a given data set in a group is determined. Then the rank order for all the sets in that group is summed giving the value of the metric. The lower the value of the metric the stronger the ranking. Table 5 shows how this is done. This table contains seven columns, one for each compression algorithm for each FOM value. The last column contains the sum of the

Table 4: Statistics Summary
Data Set A

VQ				
CR	1	8	16	32
MEAN	90.18	90.09	88.73	88.65
STD	18	16.45	13.64	12
KURTOSIS	1.42	1.74	1.92	1.99
PEAK	207.09	177	183	183
CONTRAST1	8506.67	8419.28	8097.71	8040.69
CONTRAST2	2.3	1.97	2.07	2.07
MNSD.	0.2	0.18	0.15	0.14
ENTROPY	6.16	6.03	5.71	5.02

MRES				
	1	8	16	32
MEAN	90.18	89.73	89.73	89.73
STD	18	14.82	12.99	11.71
KURTOSIS	1.42	0.7	1.11	1.23
PEAK	207.09	210.45	205.82	200.77
CONTRAST1	8506.67	8312.19	8261.24	8229.55
CONTRAST2	2.3	2.35	2.3	2.24
MNSD.	0.2	0.17	0.15	0.13
ENTROPY	6.16	5.93	5.73	5.57

JPEG							
	1	2	4	8	16	32	64
MEAN	90.23	90.18	90.19	90.18	90.19	90.01	90.02
STD	18.02	18.42	18.3	16.13	15.01	13.66	12.02
KURTOSIS	1.42	1.33	1.42	0.58	0.66	0.97	1.58
PEAK	208.1	209.4	211.5	212.95	210.6	207.25	208.4
CONTRAST1	8464.7	8469.14	8463.5	8110.79	8356.36	8313.41	8201.67
CONTRAST2	2.3	2.32	2.34	2.35	2.33	2.3	2.31
MNSD.	0.2	0.2	0.2	0.18	0.17	0.15	0.13
ENTROPY	6.16	6.2	6.22	6.05	5.94	5.6	2.63

Table 4: Statistics Summary
Data Set B

VQ				
	1	8	16	32
MEAN	84.99	86.42	84.43	84.36
STD	19.35	17.95	15.20	13.48
KURTOSIS	0.94	0.98	1.22	-0.11
PEAK	201.11	171.15	172.15	170.15
CONTRAST1	7674.93	7850.44	7478.03	7361.36
CONTRAST2	2.38	1.99	2.05	2.03
MNSD.	0.23	0.21	0.18	0.16
ENTROPY	6.27	6.16	5.89	5.46

MRES				
	1	8	16	32
MEAN	84.99	84.48	84.48	84.49
STD	19.36	16.36	14.55	13.24
KURTOSIS	0.94	0.34	0.54	0.50
PEAK	201.35	202.25	198.80	189.05
CONTRAST1	7674.93	7480.56	7424.94	7388.60
CONTRAST2	2.38	2.39	2.36	2.24
MNSD.	0.23	0.20	0.17	0.16
ENTROPY	6.27	6.05	5.89	5.75

JPEG							
	1	2	4	8	16	32	64
MEAN	84.99	84.98	84.99	84.95	84.96	84.99	84.98
STD	19.36	19.85	19.66	17.60	16.52	15.33	14.09
KURTOSIS	0.94	0.87	0.52	0.32	0.36	0.59	0.99
PEAK	201.35	202.50	207.05	210.45	203.90	199.70	198.35
CONTRAST1	7674.93	7690.25	7835.69	7598.02	7566.03	7534.41	7513.77
CONTRAST2	2.38	2.40	2.45	2.49	2.41	2.36	2.35
MNSD.	0.23	0.24	0.23	0.21	0.20	0.18	0.17
ENTROPY	6.27	6.32	6.33	6.18	6.08	5.73	2.83

Table 4: Statistics Summary
Data Set C

	VQ				
	1	8	16	32	64
MEAN	80.69	80.33	79.94	79.92	79.59
STD	21.73	19.90	17.69	16.12	14.85
KURTOSIS	1.25	1.25	2.40	2.41	3.44
PEAK	216.05	206.92	212.26	199.30	120.00
CONTRAST1	7006.78	6871.83	6725.45	6668.96	6573.54
CONTRAST2	2.55	2.35	2.45	2.29	1.51
MNSD.	0.27	0.25	0.22	0.20	0.19
ENTROPY	5.96	6.25	5.99	5.35	2.26

	MRES						
	1	2	4	8	16	32	64
MEAN	80.69	80.20	80.21	80.21	81.47	80.22	80.21
STD	21.73	21.76	20.83	18.53	17.59	16.44	15.60
KURTOSIS	1.25	1.17	0.93	1.24	1.43	1.92	2.30
PEAK	205.00	205.56	182.11	198.33	184.38	173.67	167.00
CONTRAST1	7006.78	6949.86	6890.68	6803.63	6956.58	6760.98	6701.78
CONTRAST2	2.55	2.57	2.58	2.66	2.54	2.51	2.45
MNSD.	0.27	0.27	0.26	0.23	0.22	0.21	0.19
ENTROPY	5.96	6.43	6.39	6.20	6.12	5.98	5.86

	JPEG						
	1	2	4	8	16	32	64
MEAN	80.69	80.67	80.71	80.66	80.68	80.63	80.72
STD	21.76	22.07	21.96	20.03	19.07	18.01	16.86
KURTOSIS	1.25	1.15	0.93	1.08	1.28	1.71	2.31
PEAK	203.88	206.67	207.56	214.00	204.11	201.67	197.44
CONTRAST1	7006.78	7018.70	7019.88	6930.72	6896.25	6850.34	6821.82
CONTRAST2	2.55	2.57	2.58	2.66	2.54	2.51	2.45
MNSD.	0.27	0.27	0.27	0.25	0.24	0.22	0.21
ENTROPY	5.96	6.45	6.46	6.33	6.30	5.90	2.84

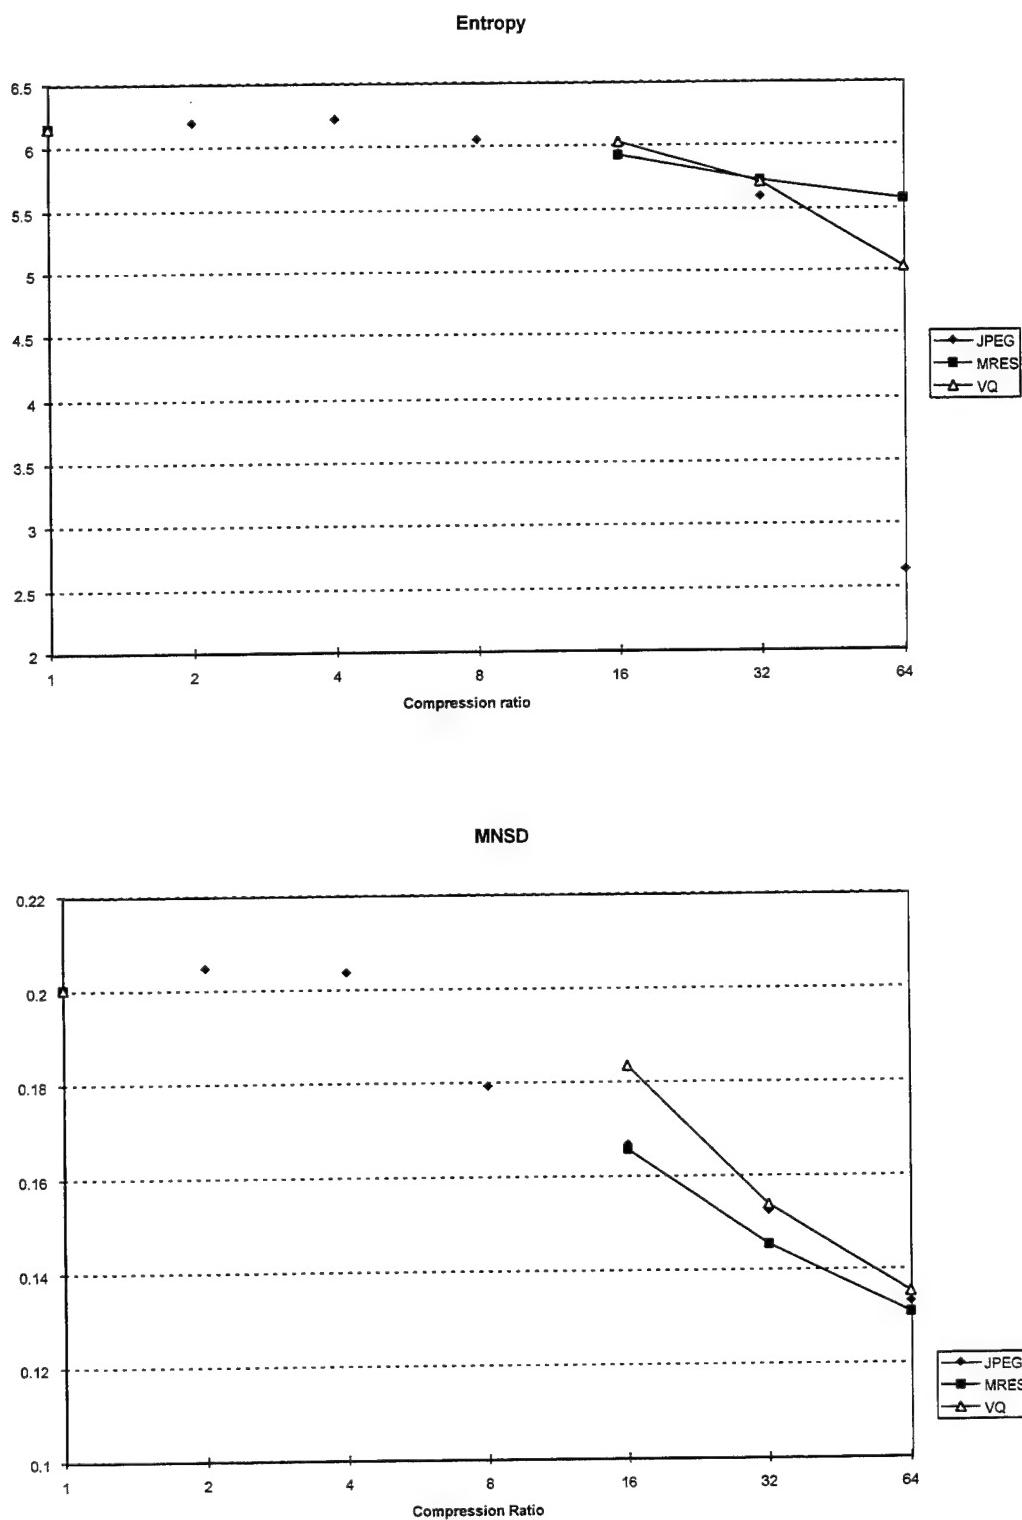


Figure 11. Entropy and MNSD as a Function of Compression Ratio

Table 5: RANKED CORRELATION STRENGTH

VALUE1

JPEG	
ENTROPY	0.93993974
MNSD	0.85699325
STD	0.85317662
CONTRAST	0.75207024
MEAN	0.72980459
PEAK	0.55959243
KURTOSIS	0.13277714
CONTRAST2	0.05256169

VALUE2

JPEG	
ENTROPY	0.96665514
MNSD	0.8448125
MEAN	0.74337636
STD	0.73572083
KURTOSIS	0.48711788
CONTRAST2	0.43895302
PEAK	0.41595093
CONTRAST	0.41118441

MRES

MRES	
ENTROPY	0.49694709
STD	0.48710198
MNSD	0.48691038
CONTRAST	0.46795628
MEAN	0.42889198
PEAK	0.40873517
CONTRAST2	0.36477729
KURTOSIS	0.25289126

MRES

MRES	
MNSD	0.87981234
ENTROPY	0.83156943
PEAK	0.82453241
STD	0.78797987
CONTRAST	0.69130259
MEAN	0.52806341
KURTOSIS	0.07125943
CONTRAST2	0.05494215

VQ

VQ	
CONTRAST2	0.93523945
PEAK	0.92221403
KURTOSIS	0.67515367
MNSD	0.49634283
STD	0.48867478
ENTROPY	0.45314177
CONTRAST	0.3730749
MEAN	0.25263304

VQ

VQ	
MNSD	0.95476131
PEAK	0.91464784
KURTOSIS	0.50285052
CONTRAST2	0.30431997
STD	0.20883093
CONTRAST	0.15648297
MEAN	0.05621667
ENTROPY	0.02641721

ranking of the previous six columns and is sorted from low value to high value i.e. from best to worst. For this particular group the best value possible is 6 meaning that the same statistic always ranked highest for each data set in the group. The lowest value possible is 48 indicating that the same statistic always ranked last.

A total of six groupings were evaluated. The results are given in Table 4. Table 4a. shows the results for all the data sets and is the same result shown in the last column of Table 3. Table 4b. and 4c. are grouped by FOM value. All three algorithms were grouped and summed for FOM values 1 and 2 respectively in b and c. The best possible value for these sets is 3 and the worst is 24. Notice that for value 2 MNSD has a value of 4, nearly unanimous. MNSD and entropy are the two best for these cases. The last three groupings in Table 4 d, 4e and 4 f show the results of grouping FOM values by algorithm. For these groups the best possible value is 2 and the worst is 16. For JPEG this value is obtained for the entropy statistic. For MRES entropy scores a 3 but is still the top ranked statistic. VQ results are somewhat different. MNSD ranks second while entropy is ranked only seventh. This last result is not consistent at all with the other results.

In summary, MNSD ranked first or second in all six groupings. Entropy finished first or second in all grouping except for VQ.

5.7 Correlation Analysis

The rank order metric shows that JPEG and MRES compression appear to behave in a similar manner e.g. the rank ordering is similar. VQ on the other hand shows significant differences. The nonparametric Spearman's Rank correlation was computed to determine the similarity of the various compression scheme behavior. The compression algorithms were paired and the resulting three correlation coefficients were computed. The results are shown in Table 5. MRES and JPEG have a high positive correlation while VQ shows a relatively weak correlation with MRES and JPEG. This is a strong indication that there is something fundamentally different about VQ.

5.8 Performance Analysis

In this section we will try to offer a possible explanation for the relatively poor performance of the VQ algorithm vice the transform algorithms. We believe the explanation is based on the relative decorrelation induced by the algorithms. In the transform algorithms the imagery is

matched to the transform kernel. By optimally matching this function to the imagery by adjusting the phase and amplitude relatively good representations of the imagery can be made. For the low frequency terms the gross shape of the image is replicated. High frequency terms replicated the detailed structure of the imagery. The process is similar to interpolation that makes the optimal fit to a predetermined curve, the curve in this case being the kernel. The superposition of all the frequency terms gives the "optimal" fit for the particular functions. Since these functions are curves they will tend to follow the gradient of the image and the resultant smoothing minimizes errors in the reconstructed image. In transform algorithms no a priori assumptions are made about the pattern of pixel values in a given image. Therefore the algorithm can determine the actual pattern on its own.

This last statement is not true for VQ. The very basis of VQ is the assumption that there is a priori knowledge of the pixel pattern in a given image. Although good estimates can be made by examining a representative sample of images there are factors that can preclude a truly representative sample is attainable. One factor that can preclude this is the calibration of the sensor. In an uncalibrated sensor even the same target will not have the same pixel values from scene to scene. Since VQ depends on having a code book of precise values this can make training extremely difficult. Our sensor is indeed an uncalibrated sensor.

Table 6: Statistics Rank Order

	JPEG-1	MRES-1	VQ-1	JPEG-2	MRES-2	VQ-2	RANK
ENTROPY	1	1	6	1	2	8	19
MNSD	2	3	4	2	1	1	13
STD	3	2	5	4	4	5	23
CONTRAST	4	4	7	8	5	6	34
MEAN	5	5	8	3	6	7	34
PEAK	6	6	2	7	3	2	26
KURTOSIS	7	8	3	5	7	3	33
CONTRAST2	8	7	1	6	8	4	34

Table 7: Rank Order by Class

OVERALL	
MNSD	13
ENTROPY	19
STD	23
PEAK	26
KURTOSIS	33
CONTRAST	34
MEAN	34
CONTRAST2	34

7a.

VALUE1	
ENTROPY	8
MNSD	9
STD	10
PEAK	14
CONTRAST	15
CONTRAST2	16
MEAN	18
KURTOSIS	18

7b

VALUE2	
MNSD	4
ENTROPY	11
PEAK	12
STD	13
KURTOSIS	15
MEAN	16
CONTRAST2	18
CONTRAST	19

7c

JPEG	
ENTROPY	2
MNSD	4
STD	7
MEAN	8
CONTRAST	12
KURTOSIS	12
PEAK	13
CONTRAST2	14

7d

MRES	
ENTROPY	3
MNSD	4
STD	6
CONTRAST	9
PEAK	9
MEAN	11
KURTOSIS	15
CONTRAST2	15

7e

VQ	
PEAK	4
MNSD	5
CONTRAST2	5
KURTOSIS	6
STD	10
CONTRAST	13
ENTROPY	14
MEAN	15

7f.

Table 8: Spearman's Rank Correlation Coefficient

MRES - JPEG	.815
JPEG - VQ	-.369
VQ - MRES	-.143

6. Entropy

In Section 5.0 we showed that the two best predictors of detection performance were entropy and MNSD. We also showed that VQ performance is in some way fundamentally different from the transform approaches. It was postulated that this difference is due to the fact that VQ tends to decorrelate the imagery, while the transform algorithms tend to maintain correlation within the imagery. Since entropy is a measure of the information content of a signal we might expect that compression algorithms that maintain the information or entropy would perform better. In fact, the following claim made for wavelets is repeated from section 2.3; "The Wavelet transform provides an efficient method of decomposing a signal into sparse and meaningful coefficients, which are highly related to the information content of the signal." In other words the wavelet transform conserves entropy.

The concept of entropy was introduced into signal processing by Claude Shannon in 1948.²⁴ In the early days of communications and information theory, a metric was needed to characterize the amount of information in a signal. Entropy, for a gray scale image, is defined as $-\sum p(x_i) * \log_2(p(x_i))$ where x_i is the range of pixel values possible in a given image (0 - 255 in an 8 bit image) and $p(x_i)$ is the relative frequency, or probability, of occurrence of each pixel value. Entropy is expressed in units of bits so that, if an entropy calculation results in a value of n, we say that the image has n bits of information. The value of entropy will range from 0 to the number of bits for which the image is constructed, commonly 8 bits in SAR imagery.

In order to get a feeling for the relationship of entropy to pixel value distributions, let's look at three examples. For example one, assume a 4 x 4 pixel, 4 bit (16 level) image. If all of the pixels have the same value, then the entropy will have entropy = 0 since the probability of this value occurring is 1, and the log of 1 = 0, regardless of the value of the pixel. For example two, at the other extreme, assume that each of the 16 pixels has a different value. Then $p(x_i) = 1/16$ and $\log_2(p(x_i)) = -4$ for each of the sixteen pixels, and the entropy of the image is 4 bits. Example three is an intermediate case where there are 8 different values and 2 pixels with each of these values, then $p(x_i)=1/8$ and $\log_2(p(x_i)) = -3$ and the entropy of the image is 6. Notice that in examples 1 and 3, the exact value of the pixels is irrelevant. The entropy will be the same for any pixel values chosen.

From the above examples, it is apparent that there is an underlying assumption in using entropy to quantify the image content of a signal. That is, that the signal must have some amount of correlation. Intuitively, digital images are not randomly valued pixels. If an image contains an object of some sort, the values of the pixels which represent that object will be related, e.g. an extended radar reflector will consist of a number of high value pixels, and the probability of a very small pixel value is very low. Likewise, clutter will consist of pixels which have a range somewhat smaller than the dynamic range of the image. The higher the correlation, the lower the entropy.

There are two other relevant characteristics of entropy. First, from the definition, it is apparent from the $-\log_2(p_i(x))$ term that the lower probability pixel values contribute most to entropy and thus information content. Intuitively, large flat areas are not very interesting in an image. However, the relatively infrequent edges contain much information about the location of objects in an image. Also, in SAR the bright return pixels make up a small percentage of the total number of pixels in an image, typically less than 1%. Second, the value of image entropy represents the number of pixels to which an image can be compressed and be restored without distortion. For example, if an 8 bit image has an entropy of 4 it, can be theoretically compressed by a factor of 2, and the exact original image can be recovered.

SAR imagery is characterized by coherent speckle²⁵. This superimposes a random component to the signal, which tends to decorrelate the SAR image. As a result, one would expect the entropy of SAR imagery to be rather high implying a high information content. As discussed above, high entropy is not the same as useful information if the signal is random. Therefore, information content of SAR imagery which can be used by a human interpreter or an automatic detection or recognition algorithm may be lower than the entropy would imply. However, it would be expected that the entropy is related monotonically to image information content. As with other types of imagery, it would be expected that higher resolution imagery would have lower average entropies than lower resolution imagery. Higher resolution imagery has more redundant information, i.e., it is more correlated.

In view of these considerations Table 3 in section 5.6 was recomputed using only the MRES and JPEG data. The results are shown in Table 6.

Table 6: Rank Order

	JPEG-1	MRES-1	JPEG-2	MRES-2	RANK
ENTROPY	1	1	1	2	5
MNSD	2	3	2	1	8
STD	3	2	4	4	13
MEAN	5	5	3	6	19
CONTRAST	4	4	8	5	21
PEAK	6	6	7	3	22
KURTOSIS	7	8	5	7	27
CONTRAST2	8	7	6	8	29

As the table shows entropy is the overall best predictor and is best in 3 of 4 cases. In the fourth case entropy is second. Contrast this with Table 3 where MNSD is best overall with entropy finishing as low as seventh of eight for VQ overall

These observations in the previous paragraph were quantized by calculating the correlation coefficient relating compression levels and the FOM for each case. These results are shown in Table 2. These results bear out the subjective observations. For JPEG, the correlation is very high for both values of the FOM. This would indicate that entropy is useful for prediction detection performance based on entropy measurements. MRES is more strongly correlated than was expected, and entropy's value of detection performance is not as good, but it appears as if it would still be useful, especially at performance levels corresponding to FOM parameter 2. The VQ are essentially uncorrelated and of no use in predicting detection performance. No explanation has been found for this behavior, although it should be pointed out that both JPEG and MRES are based on transforms of the original image data. These results also indicate that FOM parameter value 2 yields better correlation than value 1. This can probably be attributed to the fact that value 2 represents a more "stable" area of the ROC curves. VQ again is not consistent with this observation, indicating that its behavior is random in nature.

Table 7: Correlation Coefficients

FOM PARAMETER	MRES	JPEG	VQ
Value 1	-.49695	-.93994	-.45569
Value 2	-.83157	-.96666	-.02416

7. Conclusions

7.1 Prescreener Performance

It is clear from the results that JPEG is superior to the VQ algorithm used in this study. This holds true for all 3 data sets. JPEG is also superior to MRES at compression rates up to .5 bpp and about the same at .25 bpp. However, at .125 bpp, MRES appears to be the preferred algorithm. JPEG is the algorithm of choice for rates up to .25 bpp. Its performance is equal to or better than the other algorithms, and is a widely used standard. At very high compression rates, the MRES would appear to be the best algorithm. However, this assumes that the entropy of SAR imagery actually are an indication of the utility of the imagery.

In general, the transform based algorithms performed better and in particular were better "behaved". Vector quantization performance was erratic, but in some cases provided superior performance. The possibility exists that more sophisticated code book development might result in an algorithm which exhibits superior performance. However, because VQ appears to be more dependent on image content several code books may have to be developed to deal with imagery collected in various situations.

Another conclusion from this study is that compression algorithms and ATR algorithms need to be matched to each other. This is apparent from looking at the degradation differences between the prescreener with and without the fractal dimensionality discriminant. Compression had a significantly greater affect on the discriminator version of the algorithm. It is possible that certain compression algorithms might be part of the ATR algorithm suite serving as detection algorithms. A related issue is that the prescreener was not trained on compressed imagery. Had training been done on compressed imagery separately for each compression algorithm, as opposed to uncompressed imagery, degradation may not have been as great.

Given the above observations, it can be said that it appears that the prescreener reported on in this paper can provide satisfactory performance robustness for meaningful compression ratios, e. g. 8:1. It is our judgment that compression ratios of 8:1 and higher are required to meet the communications needs of future systems.

The prescreener evaluated here is only the first stage of a multistage algorithm. One of the goals in tuning the prescreener is to ensure that a minimal number of targets are missed at this stage. As a result a higher false alarm rate can be tolerated than if the prescreener output were the final product. Additional false alarm reduction occurs in later stages of the algorithm suite.

Since the amount of area covered in this study is somewhat limited the statistical validity of the results is somewhat marginal. Additional studies using significant areas should be done before selecting an algorithm for operational use. However, the results presented here indicate that compression of SAR imagery can be done with minimal impact to automatic target detection algorithm's performance at useful compression ratios.

7.2 Statistics

In general, image statistics are a poor indicator of automatic target detection performance. Of the eight statistics examined in this paper only the MNSD and entropy metrics show consistently high correlation to detection performance. The target screening experiment showed that for the JPEG and MRES, entropy does appear to be a rough indicator of image utility. Even in the case of these two metrics the VQ results showed inconsistencies with the other algorithm results and affects detection significantly. Non-parametric analysis of the paired algorithm results indicate that VQ behaves differently than transform based algorithms.

8. Recommendations for Further Study

There are several areas of further study indicated by the results shown here. First, MRES, or another wavelet-based algorithm, should be used to compress data sets A and B at .125 bpp to see if the results for set C are repeated for this data. Compression ratios should then be increased to see if entropy values remain high for extremely high compression. Secondly, additional experiments should be done to validate the correlation between entropy and target screening for transform-based algorithms. These experiments should include data sets B and C. Other automatic detection/recognition algorithms should also be evaluated. The reason for VQ entropy not correlating to target screening performance should also be investigated. A third area of interest would be to determine image analyst performance as a function of compression ratio and algorithm. Penrod and Kuperman²⁶ have investigated this for DCT, Block Truncation, and a different VQ implementation, and found little affect on analyst performance. However, the compression rate was only 2 bpp. It would be useful if entropy could be used to predict analyst or automatic target recognition performance based on the entropy of the image.

ACKNOWLEDGMENT

We would like to thank Dr. Fred Garber of Wright State University for his encouragement, advice, and for his assistance in reviewing this report.

9.0 References

- [1] Wharton, J.J., Gorman, J. D., Werness, S. A., Gleason, J. M., ARDA Interim Report, March 1993
- [2] Wharton, J.J., Gorman, J. D., Werness, S. A., Gleason, J. M., Ibid.
- [3] Wharton, J.J., Gorman, J. D., Werness, S. A., Gleason, J. M., Ibid.
- [4] Wharton, J.J., Gorman, J. D., Werness, S. A., Gleason, J. M., Ibid.
- [5] G. Wallace, "The JPEG Still Picture Compression Standard", Commun. ACM, Vol 34, pp 33 - 45, 1991.
- [6] MIL-STD-188-198A, 15 Dec 1993, Joint Photographic Experts Group (JPEG) Image Compression for the National Image Transmission Format Standard.
- [7] A. Gersho and R. Gray, "Vector Quantization and Signal Compression", Kluwer Academic Publishers, 1992.
- [8] P.A. Chow , T. Lookabaugh, and R. M. Gray, "Optimal Pruning with Applications to Tree Structured Source Coding," IEEE Trans. Inform. Theory, Vol 35, pp 299 - 315, 1989
- [9] Wharton, J.J., Gorman, J. D., Werness, S. A., Gleason, J. M., Op. Cit.
- [10] R. M. Gray and J. Linde, "Vector Quantization and Predictive Coding Quantization for Gaussian Markov Sources," IEEE Trans. Comm., vol COM - 30, No. 2, pp 381 - 389, February 1982.
- [11] Fastman Inc. "The Wavelet Handbook". Technical Report, Defense Advanced Research Projects Agency, 1991 (AD-B151 677).
- [12] Fastman Inc., "The Wavelet Handbook". Ibid., page 2.
- [13] Wharton, J.J., Gorman, J. D., Werness, S. A., Gleason, J. M., Op. Cit.
- [14] I. Daubechies, "Orthonormal Basis of Compactly Supported Wavelets", Communications Pure and Appl. Math, Vol 44, pp 906 - 996, 1988.
- [15] M. G. Perkins and T. Lookabaugh, " A Psychophysically Justified Bit Allocation Algorithm for Subband Image Coding Systems," Proc. ICASSP, pp 1815 - 1818, 1989.
- [16] Lewis, April A., "Chicken Little Grayling 1992 Final Ground Truth Data Report", April 1993.
- [17] W. K. Pratt, "Digital Image Processing", John Wiley and Sons, Inc, 1991.
- [18] Raney, S. D., Nowicki, A. R., Record, J. N. and Justice, M., "ARAGTAP ATR System Overview", Proceeding of The 1993 National Fire Control Symposium.
- [19] W. K. Pratt, "Digital Image ", John Wiley and Sons, Inc, 1991.

- [20] Tom, Victor T., Peli, Tamar, and Keydel, Eric R., "Morphology -Based Target Detection and Discrimination," Proceedings, "Detection, Discrimination & Classification of Targets in Clutter," Huntsville, Alabama, 13-15 November 1990.
- [21] Peli, T., Tom, V., Lee, B., "Multi-Scale Fractal and Correlation Signatures for Image Screening and Natural Clutter Suppression, "SPIE Vol. 1199 Visual Communications and Image Processing IV (1989).
- [22] Peli, T. , "Multiscale Fractal Theory and Object Characterization," Vol. 7, No. 6, Journal of the Optical Society of America A.
- [23] The Khoros Group, Khoros Release 1.0, Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM, Online Documentation.
- [24] Shannon, C.E., "A Mathematical Theory of Communications", Bell Systems Technical Journal, Vol 27,pp 379-423, July 1948, pp 623 -656, Aug 1948.
- [25] Morris, Guy V., "Airborne Pulsed Doppler Radar", Artech House Inc, Norwood, MA 1988, pp 183-184.
- [26] Penrod, Todd D. and Kuperman G., "Image Quality Analysis of Compressed Synthetic Aperture Radar, Air Force Armstrong Laboratory Technical Report, January 1993.